Poster 27: Absolute Localization for Surface Robotics in GPS-Denied Environments using a Neural Network

A.S. Chung*[1], P. Ludivig*[2,3], R.W.K. Potter*[4,5], T. Seabrook*[6], and B. Wu*[7,8]

*Equal Contribution, [1]Tensorlicious, [2]ispace Europe, [3]University of Luxembourg, [4]GTA Travel, [5]Brown University, [6]University of Oxford, [7]NVIDIA, [8]National Astronomical Observatory of Japan

1. Introduction

- Accurate localization in planetary surface robotics is essential for navigation, path planning, and science objectives.
- On Earth, absolute localization can be readily achieved via satellite navigation (e.g., GPS). However, such systems are unavailable on other planetary bodies such as the Moon or Mars.
- Current methods rely on time- and labor-intensive human visual matching of surface perspective features with satellite images. Relative localization via visual and inertial odometry accumulate errors over time and lead to inconsistencies.
- Thus, a method that can quickly, automatically, and accurately reduce the position search space is of great benefit to future planetary exploration missions. **This project[1] presents a new approach to localizing planetary rovers: training an artificial neural network to match surface-perspective imagery to corresponding satellite maps.**

2. Methods

Simulated Environment:

- A simulated environment was used to generate a dataset adequate in size for training a deep neural network.
- The synthetic Lunar surface environment was assembled in Unreal Engine 4[2] using MoonLandscape v3.0[3].
- Distinct zones were set for **training**, **validating**, and **testing**.

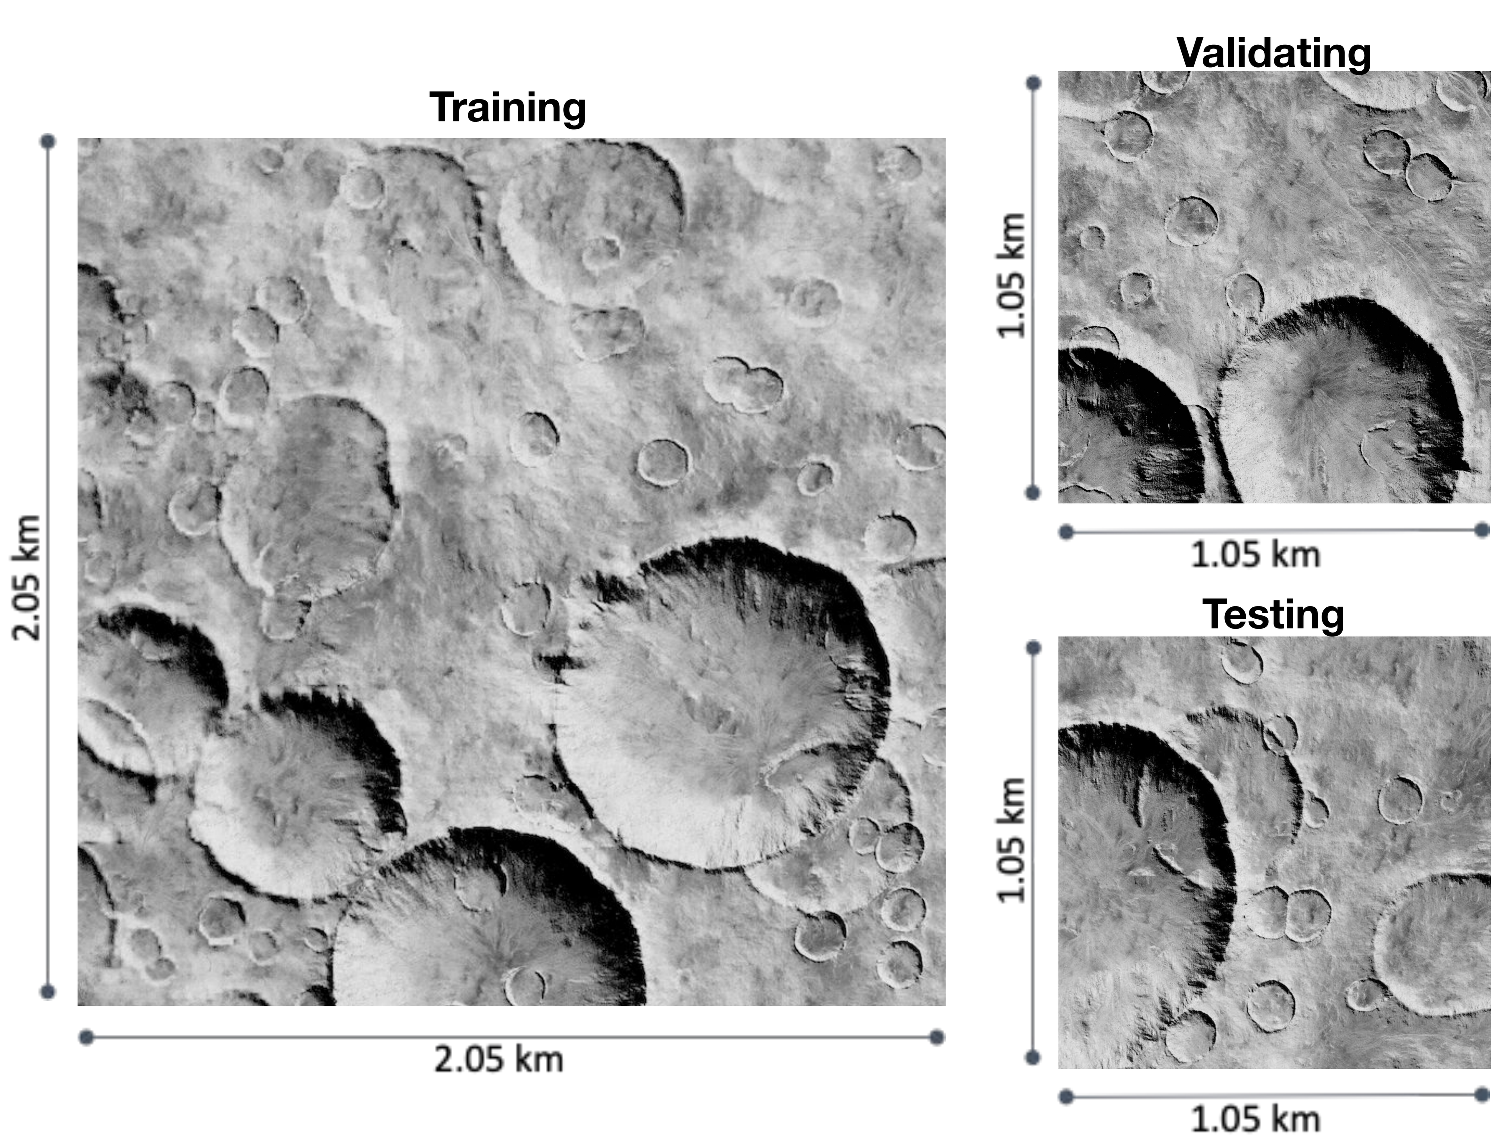

3. Methods (cont'd)

Dataset Generation:

- A virtual rover camera was spawned at distinct random locations (N~600,000) throughout the simulated environment, capturing surface-perspective images in the 4 cardinal directions as well as the corresponding ground truth orbital map.

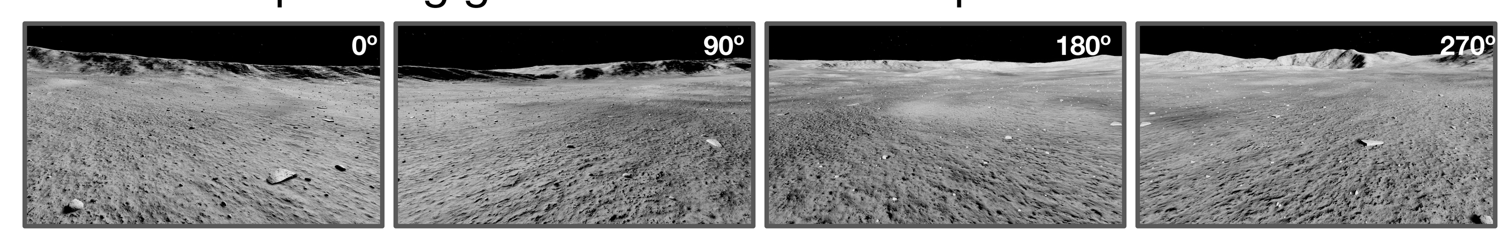

Dataset Processing:

- Each set of 4 surface perspective images was then reprojected into an approximate aerial view using rover camera properties and assuming that the terrain is locally flat.

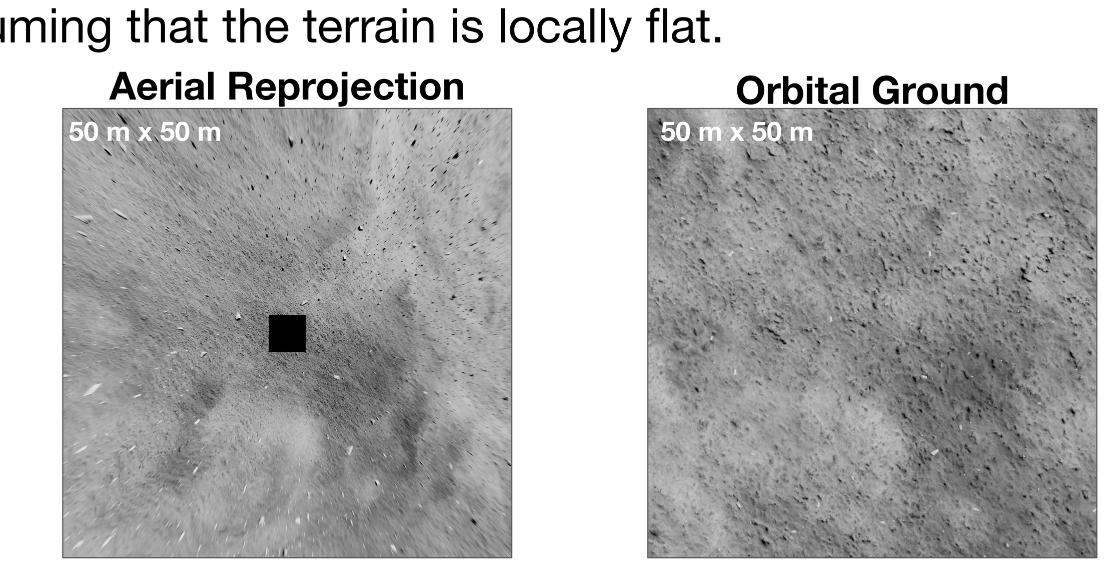

- The reprojection is paired with an orbital map of the same location (*matching*) or one from a different (*non-matching*) location with 50/50 probability, downsampled, and fed into the neural network.

Neural Network:

- ***PLaNNet*** (**P**lanetary **L**ocaliz**a**tion **N**eural **Net**work), a Siamese neural network, was trained to classify pairs of reprojections and orbital maps as *matching* or *non-matching*:
  - Each image enters a pre-trained ResNet-50 feature extractor
  - The feature vectors are concatenated, fed into a 256-neuron fully connected layer (30% dropout) to produce the final match/no-match logits vector, and softmax is applied to produce the match/no-match probability distribution for a pair of inputs.

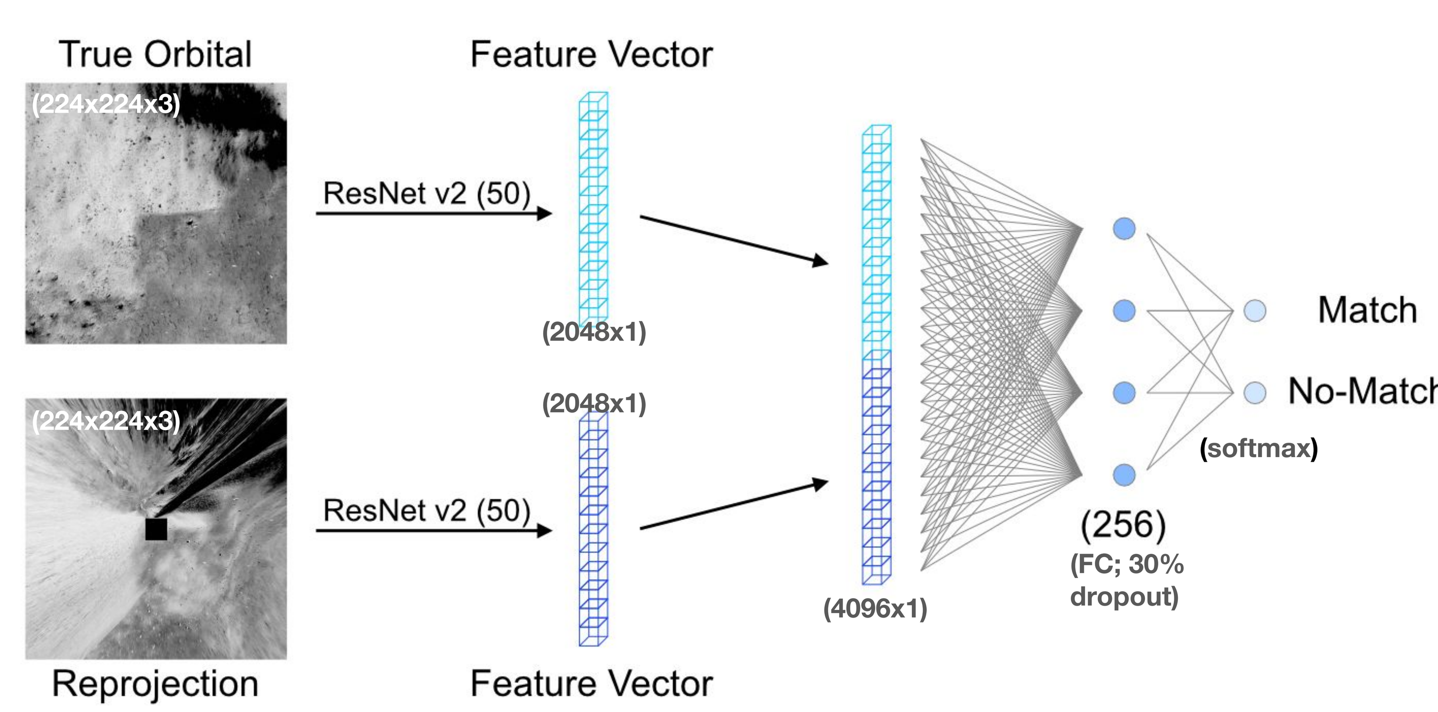

4. Results

Public Dataset and Code:

- We produced a publicly available synthetic Lunar dataset and open source code for processing, training, and benchmarking localization algorithms[4]. The dataset contains 2.4+ million surface-perspective images at 600,000+ distinct locations split among the training, validating, and testing zones.

Image Type	Number	Physical Scale	Original Resolution	Downsampled Resolution
Surface Perspective	2.42×10^6	90° x 50.6° FoV	1920 x 1080 px	-
Reprojection	6.06×10^5	50 m x 50 m	1000 x 1000 px	224 x 224 px
Orbital Ground Truth	6.06×10^5	50 m x 50 m	1000 x 1000 px	224 x 224 px

Localization:

- The reprojection is compared against an array of candidate locations via a sliding window over any given orbital map. *PLaNNet* calculates probabilities of a *match* with each candidate.

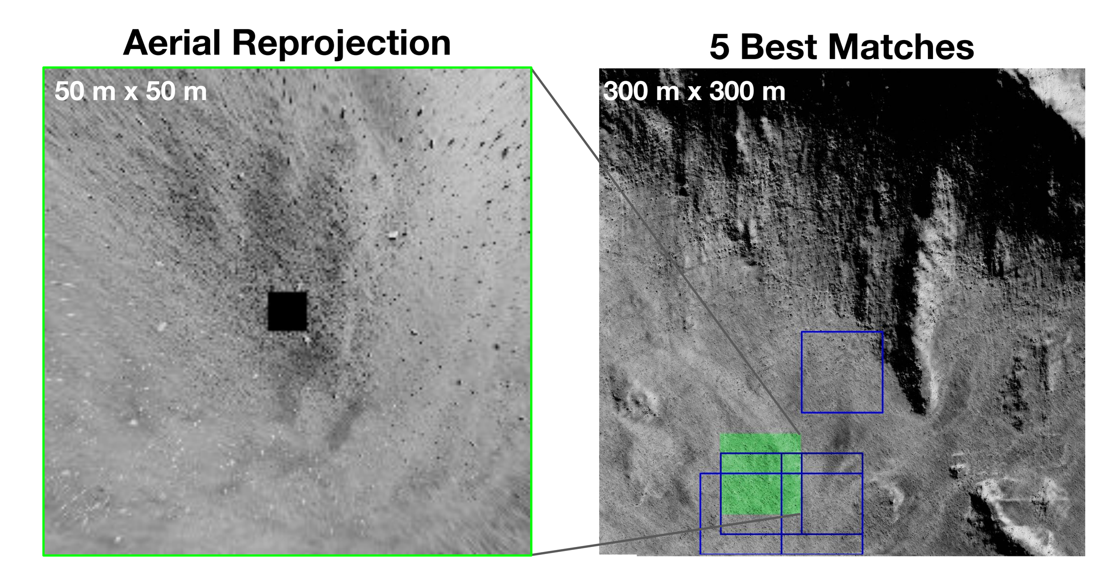

- Experiments using random locations within testing zone:
  - 50 locations in $(300\text{m})^2$ subregion (3600 candidates)
  - 300 locations in full $(1.05\text{km})^2$ testing zone (40401 candidates)
- In general, *PLaNNet* returns a location within 5m of ground truth from the top 10% inferences from available candidate regions (i.e., **90% reduction of search space**). It performs >2x better than standard computer vision benchmarks (SAD/SSD/random).

Conclusions: This proof-of-concept demonstrates promising capabilities for neural network approaches to absolute localization in remote planetary surface environments. Work is in progress to include stereo camera depth information and new architectures.

Acknowledgments:
This work was published at IROS 2019[1] and originated from the NASA Frontier Development Lab (2018). Support was provided by the SETI Institute and NASA Ames Research Center. Funding was provided by Intel, Space Resources Luxembourg, and XPRIZE; compute was provided by Intel. PL thanks the Luxembourg National Research Fund (FNR 11824057).

References:
[1] Wu, B. et al. (2019) *IEEE IROS*, 3262. [2] www.unrealengine.com [3] www.unrealengine.com/marketplace/the-moon [4] http://moonbench.space/